CHAPTER 6

VARIABILITY OF BENEFIT-COST ESTIMATES

I. Risk and Uncertainty

Previous chapters discuss various aspects of economic analysis. Such analysis is almost always characterized by uncertainty in that it involves the use of estimates, forecasts and assumptions related to key variables. For example, the benefits or costs of a given undertaking are typically not known with certainty but must be estimated. In addition, it is often the case that some benefit or cost value is estimated for a single year or other relevant unit of time and then is projected to grow at some rate (which itself may change overtime) out into the future. Or, as discussed in Chapter 5, the relevant evaluation period may also be subject to uncertainty.

This chapter considers how various types of uncertainty can be characterized and how they impact cost-benefit analysis. First, it is important to differentiate between "uncertainty" and "risk." In those cases when it is possible to characterize the uncertainty numerically or mathematically, one can analyze the potential variation in benefit-cost analysis results. Such analysis, sometimes termed “risk analysis,” is the focus of this chapter. It is important to emphasize that such analysis requires one to be able to characterize uncertainty to the maximum degree of mathematical specificity possible.[1] The type of analysis of the variability of benefit-cost estimates (as discussed below) is influenced by the degree of mathematical specificity.

In order to fully inform decision-makers of the consequences that may, but are not guaranteed to follow from their actions as a result of uncertainty, analysis of the risk of variability of estimates should be conducted as part of every benefit-cost analysis. This analysis should, at a minimum, report expected value estimates for benefits, costs, and net present values; identify the key sources of uncertainty; and estimate the impact of these uncertainty sources on outcomes. This may be accomplished by undertaking a sensitivity analysis. Where relevant, probability distributions of benefits, costs, and net present value should also be presented. Stochastic simulation and other methods can be used to derive these results. This chapter presents methods for accomplishing these requirements as well as other techniques useful for analyzing risk.

---

[1] The term “risk analysis” as used in this chapter pertains to the evaluation of the variability associated with benefit-cost estimates. The term “risk analysis” or “ risk assessment” is also used to refer to various techniques used to estimate accident risks when historical data are absent or limited. A summary of some of these techniques is provided in Chapter 3, Section III.A.4.

II. Certainty Equivalents

An understanding of the concept of certainty equivalents is important to any risk analysis. A certainty equivalent refers to the net benefits of a *certain* (i.e., non-risky) return that has the same value to individuals as the expected value of an uncertain (i.e., risky) return. For example, suppose the expected present value of an undertaking that entails risk is $1000. If individuals would be willing to trade these risky net benefits for a certain lump sum of $800, then $800 is the value of the "certainty equivalent" of the risky activity. When the certainty equivalent is less than the risky expected value of an activity, risk aversion is implied (i.e., there is a preference to avoid risk). The difference between the expected value and the certainty equivalent value is called the "risk premium." When the certainty equivalent is equal to the expected value of the undertaking, risk neutrality is implied.

For most FAA benefit-cost analyses, expected values should be treated as identical to certainty equivalent values. This is because, to the extent that costs and benefits accrue to the Federal government, the government is large enough to be considered risk neutral with respect to FAA projects.[2] For private parties, the incremental impacts of FAA investments and regulations associated with uncertainty typically alter gains or losses over a fairly small range causing analysts to frequently treat private parties subject to FAA regulation as if they are risk neutral over this range.[3] Moreover, benefits and costs of FAA investments and regulations actually realized can be expected to approximate expected values. FAA investments and regulations are diversified over a large number of activities such that variations in benefits and costs can be expected to actually average out. Also, their impact tends to vary independently across numerous individual parties such that variations should also average out.[4]

In general, risk should be captured through the use of certainty equivalent values. For most FAA benefit-cost analyses, expected values should be used as certainty equivalents. In those rare cases where risk neutrality is not an appropriate assumption, any allowance for uncertainty should be made by adjusting the corresponding benefit or cost expected values so that they represent

---

[2] Office of Management and Budget (OMB) Memorandum, “Economic Analysis of Federal Regulations Under Executive Order 12866,” p. 21.

[3] There is a body of opinion and experience that argues that the general public is risk adverse with respect to aviation safety. However, the difference between “certainty equivalent” value and “expected” value of events has not be adequately quantified.

[4] Boardman, et. al, *Cost-Benefit Analysis: Concepts and Practice*. Prentice Hall: Upper Saddle River, NJ., 1996, pp. 225-228.

certainty equivalents.[5] It should be noted that variations in the discount rate are not an appropriate method of adjusting expected values for the risks of a particular undertaking.[6]

The remainder of this chapter focuses on risk assessment (identification of different types of risk and characterizing them mathematically) and two conventional methods of conducting risk analysis (sensitivity tests and Monte Carlo analysis). A somewhat less popular but also useful tool — decision analysis — is also discussed.

III. Risk Assessment of Benefit-Cost Results

This section describes the components of a risk assessment and discusses how the information from such an assessment can be collected. This information is then used as input to the risk analysis methodologies described in later sections. There will be focus on the following topics:

- Identification of types of risk relevant for FAA
- Discrete risky events vs. risk continuums
- Evaluating and characterizing risk probabilities
- Identifying and evaluating risk severity and impacts for discrete events
- Interdependencies among risks

Ideally, the consideration of all of these topics together would be completed in order to conduct a complete risk assessment; such an assessment would strive to identify and quantify all relevant risks for the project under evaluation. In practice, it may be difficult to explicitly consider all possible sources of uncertainty; in the simplest case, one might want to focus on just one or a few aspects of uncertainty that can be treated through sensitivity tests. This approach is considered in Section IV below. More formal risk analysis approaches that require a more complete risk assessment involving some or all of the topics listed above are considered in Sections V and VI.

In presenting the topics, it will be useful to use a simplified example as a common theme where all of the different aspects of a risk assessment can be presented. The example is that of a promising ground traffic management safety technology for aircraft. To keep the analysis straightforward, we will assume the following:

---

[5] For a discussion of the problems associated with developing such adjustments, see Boardman, et. al, *Cost-Benefit Analysis: Concepts and Practice*. Prentice Hall: Upper Saddle River, NJ., 1996, pp. 213-225.

[6] “OMB Circular A-94” (Revised--October 29, 1992) p. 13, and Edward M. Gramlich, *A Guide to Benefit-Cost Analysis--Second Edition*, Prentice Hall, Englewood Cliffs, NJ, 1990, pp. 90-100.

- Basic research begins in Year 1 (today) and is expected to take one year. (The research is assumed to be successful. This assumption will be relaxed in Section V.)

- Development and testing begins in Year 2 and is expected to take one additional year. (Development and testing is assumed to be successful. This assumption is also relaxed in Section V.)

- The technology will then be installed at the rate of one machine per 100,000 operations at each of two airports in Year 3. Each machine costs $250,000. Airport operations are assumed to grow at 3% annually from current levels for the foreseeable future.

- The technology is expected to reduce the likelihood of an accident (or collision) on the ground between two aircraft by 50 percent.

- Once installed, the technology is expected to have a useful economic life of 5 years, and to cost $100,000 per year to operate.

A sample benefit-cost flow analysis based on this information is shown in Table 6-1; the analysis is presented using constant dollars.[7]

A. Risk Types

There are a variety of different types of risk that may affect FAA investment projects. Although the following discussion may not be exhaustive, it should give some guidance to the analyst attempting to identify risk characteristics that may need to be included in a risk assessment.

For many FAA projects, an obvious area of uncertainty is the degree to which a project or expenditure may affect safety or accident rates. Using the above example, suppose the current probability of an accident is 1 per 10,000,000 operations; the new technology is expected to reduce this by 50 percent. It is the uncertainty associated with this latter number that may need to be addressed in a risk analysis. One must be able to characterize this uncertainty mathematically in order to complete a standard risk analysis.

---

[7] For ease of explanation, the analysis covers only seven years, which covers the life cycle of the first set of installed equipment. Note, however, that the analysis should be extended beyond this period if the technology would still potentially produce benefits.

TABLE 6-1

BENEFIT-COST FLOW ANALYSIS for GROUND TRAFFIC MANAGEMENT TECHNOLOGY ($000)

	Year 1	2	3	4	5	6	7
Research Cost	(1,000)						
Development and Testing Cost		(500)					
Airport 1							
Operations @3% annual growth	370,000	381,100	392,533	404,309	416,438	428,931	441,799
Number of Machines Required			3	4	4	4	4
Incremental Units Required			3	1	0	0	0
Airport 2							
Operations @3% annual growth	180,000	185,400	190,962	196,691	202,592	208,669	214,929
Number of Machines Required			1	1	2	2	2
Incremental Units Required			1	0	1	0	0
Total Incremental Machines Required			4	1	1	0	0
Cost per Machine (F&E)			(250)	(250)	(250)	(250)	(250)
Total Annual Investment			(1,000)	(250)	(250)	0	0
Total Machines Operating			4	5	6	6	6
O&M Cost per machine-year			(100)	(100)	(100)	(100)	(100)
Total O&M Costs			(400)	(500)	(600)	(600)	(600)
Total Cost	(1,000)	(500)	(1,400)	(750)	(850)	(600)	(600)
Discounted Cost Flow @7%	(1,000)	(467)	(1,223)	(612)	(648)	(428)	(400)
Accident Rate per 10,000,000 Operations without Technology			1	1	1	1	1
Reduction in Accident Rate with Technology			50%	50%	50%	50%	50%
Expected Number of Accidents							
Without Technology			.058	.060	.062	.064	.066
With Technology			.029	.030	.031	.032	.033
Cost per Accident			(70,000)	(70,000)	(70,000)	(70,000)	(70,000)
Accident Benefit with Technology			2,042	2,103	2,167	2,232	2,299
Discounted Benefit Flow @7%			1,784	1,717	1,653	1,591	1,532
Net Disc. Benefit-Cost Flow @7%	(1,000)	(467)	561	1,105	1,004	1,163	1,132
Net Present Value	3,498						
Benefit/Cost Ratio	1.73						

A second source of uncertainty relates to technical or R&D risk; this is likely to be relevant in situations where a project is being evaluated before some or all of the required research and development is completed. Thus, there may be uncertainty as to what level of effort is required and/or the operational characteristics of the technology once it is developed. In our example, any benefits from the project depend entirely on a successful research phase, so one must be able to characterize the likelihood or degree of success of the research phase if a risk analysis on this aspect of uncertainty is to be carried out.

A related uncertainty is schedule risk — the uncertainty associated with when a project may be completed or deployed. In practice, there may be a number of so-called "decision points" in the project undertaking when key decisions about whether or how to proceed can be made. In the above example, one such decision point (not reflected in the benefit-cost flow analysis) might occur at the end of Year 1 when a decision about whether to continue on with development and testing work can be made. Accurate identification of decision points can be a key part of risk analysis because the decision today to proceed on a project may well depend on outcomes that are not known today but will become known in the future.

Uncertainty with respect to implementation or compliance may also be important in the context of certain types of FAA investment projects. In our example, the technology is assumed to be installed in Year 4. But due to financial, political or other constraints, there may well be uncertainty with regard to how fast the technology can be deployed. In any event, it may be worth assessing the risk arising from departures from expected deployment of the technology. In practice, such risks may take the form of changes to the expected time path of various cost or benefits flows.

Finally, there may be risk due to uncertain estimates of future external activities or measures that affect future costs or benefits. For example, the overall dollar benefits of the ground traffic management technology will depend not only on accident rates, but also the overall number of operations at the airports in question. As noted earlier, suppose it is assumed that over the life of the technology, airport operations are assumed to grow at 3% annually. Uncertainty in this estimate may affect the overall cost-benefit calculations, and the analyst may want to attempt to address this during the risk analysis.

This discussion highlights the general point that, in principle, *any* of the elements of the cost or benefit streams in a project evaluation may be subject to uncertainty, and a good risk analysis should attempt to identify those likely to have an important impact on the results of the evaluation.

B. Characterizing Risk: Discrete Events vs. Risk Continuums

In creating a net present value benefit-cost flow analysis as described in Chapter 5, the underpinnings of the analysis will be made up of certain variables whose values affect the benefit-cost flow; in the ground traffic management technology example, one of these variables would be the projected growth rate of airport operations. An important practical aspect of risk analysis is whether the uncertainty of relevant variables can be characterized in terms of well-defined alternative discrete events; if so, and if the analyst can also assign probabilities and explicit outcomes to each possible event, then the risk analysis can be cast in a so-called "decision analysis" framework. For example, suppose the analyst can assume that the operations growth rate will be 3% with probability 0.50, and 2% or 4% each with probability 0.25. If this discrete assignment of outcomes were possible for all such variables of interest (along with their associated impacts on costs and benefits), then a decision analysis could be undertaken as described in Section VII below.

This process of identifying all possible outcomes and their likelihood of occurrence is equivalent to specifying the "probability distribution" for each variable. Of course, the probabilities across all outcomes for a given variable must add up to one. When all possible outcomes are explicitly identified as in the above example, one is essentially employing a "multinomial" distribution to characterize the uncertainty in the variable of interest.

In practice, certain types of risk identified above are more likely to be amenable to the discrete sort of analysis described above than others. For example, the uncertainty associated with accident risks often can be characterized by just a few discrete possibilities, e.g., either the technology will work as projected, leading to a 50% reduction in accidents, or it will not work at all, leaving the accident rate unchanged. In addition, if an accident does occur, the cost will be some known amount.

Also, it is often useful to characterize uncertainty associated with schedule and/or implementation risks by identifying critical decision points where Yes/No decisions are posited at some point in the future; the key here is to structure the analysis so that the decision points come immediately *after* some new piece of information comes to be known. Our example from above is set up so that a decision about development and testing can be made after one knows the outcome from the research phase. In this way, the impact of uncertainty about the research schedule itself (e.g., it may take more than one year) is mitigated because a decision to continue the project is delayed until after the outcome from the research is known, at which point a reassessment of future benefit and cost flows can be made. It is important to understand that allowance for a post-research decision point helps to decide *today* whether to undertake the project at all. How this plays itself out in a decision analysis framework is discussed in Section VI below.

Technical and other risks may or may not be adequately captured by identification of discrete outcomes. As noted earlier, there may be uncertainty as to what level of effort is required and/or the operational characteristics of a technology once it is developed. In some instances, it may be reasonable to confine the analysis to just a few possible discrete outcomes (e.g., the level of effort will be either 10 man-years, 50 man-years or 100 man-years), but in other cases it may not (e.g., the 10 and 100 man-year estimates represent low and high estimates, but anything in between — or even outside the likely range — is also possible). In the latter case, the analyst still may be able to formalize the uncertainty without explicitly identifying all possible outcomes.

A common way to formalize uncertainty is to employ a probability distribution that requires the analyst to specify only a small number of parameters, even though a whole range of outcomes is possible. For example, it is very common to use the normal distribution, which is defined by its mean (the average outcome) and variance (the spread of outcomes away from the mean). The normal distribution is a reasonable choice when the uncertainty can be characterized as a bell-shaped curve, with the most likely outcome in the middle of other outcomes (e.g., 50 man-years might be the mean outcome in the above example) and with the uncertainty spread symmetrically around the mean (e.g., if 50 is the mean (most likely) level of effort, then the probability that the level of effort might in fact turn out to be 70 is equal to the probability that it might in fact turn out to be 30).

Again, the analyst need not specify each and every possible outcome; rather, there are just two parameters that must be specified in order to completely define the normal distribution — the mean and the variance; the square root of the variance is known as the "standard deviation". As should be clear from above, the mean should be assigned the most likely (or "expected") value for the variable in question. Estimates of the appropriate variance to specify can be aided by knowing that about 95% of the probability for any normal distribution lies within 2 standard deviations of the mean. Thus, for a variable assigned a normal distribution of, say, mean 50 and variance 25 (implying a standard deviation of 5), the most likely outcome is 50, and this implies a 95 percent chance that the variable value will fall between 50±10, i.e., between 40 and 60, and only a 5 percent chance that the variable will fall outside this range.

Of course, there are circumstances when the normal distribution may not be a reasonable representation of the uncertainty associated with a variable; many other probability distributions can be specified, including the Poisson (often appropriate for characterizing accident or other events that occur infrequently), uniform (appropriate when a range of values are equally likely) and the exponential (appropriate when there is uncertainty over the length of time between certain events occurring). For practical purposes, the triangular distribution is commonly used; this distribution is characterized by a single most likely value and minimum and maximum values, with the probabilities declining linearly from the most likely to the minimum and maximum values. Details on these and other distribution types can be found in most business statistics textbooks. As described in Section V below, there is also commercial software available that helps lead the analyst through the process of selecting distributions and calculating results.

In practice, the analyst's judgment will have to be used in assessing whether it is reasonable to specify a small set of outcomes to represent all possibilities for the variables identified for consideration in the risk analysis; as noted earlier, a complete listing of explicit outcomes and associated probabilities would be required in order to employ the decision analysis approach described in Section VI below. If it is more reasonable to employ probability distributions that statistically account for uncertainty without having to explicitly identify each possible outcome, then the Monte Carlo approach to risk analysis can be used. This approach is discussed in Section V below.

C. Qualitative and Quantitative Risk Estimates: Prior and Posterior Probabilities

As hinted at above, the risk analyst may be required to make many judgments about how to characterize the uncertainty inherent in any project evaluation. In practice, much information may have to be gathered from decision makers, technical personnel, and others who have specialized knowledge about the likely risks and uncertainties. Yet it is to be expected that these sources may not be able to provide the needed information in the strict form of probability distributions or a complete set of discrete event probabilities. Instead, much of the information may be qualitative in nature (e.g., "we expect the level of effort to be 50 man-years, but it could go much higher if we have to use Approach B instead of Approach A; the chances that A will not work are fairly low, however"). Obviously, this sort of information does not by itself yield a probability distribution; rather, it is the job of the risk analyst, in concert with the personnel supplying the information and others, to make reasonable judgments in order to translate the qualitative information into quantitative information that can be used in a formal risk analysis.

Another possibility when gathering information is that the experts involved may provide data or estimates that are to be used as "second opinions." In other words, the analyst may already have a "prior" estimate of the value of some variable (based on current information), but another opinion from an expert is acquired. The expert's estimate of the value of the relevant variable may well differ from the prior estimate. If the analyst also has some information about the likely reliability of the expert's opinion, then the new "sample" information from the expert can be used to form a so-called "posterior" estimate; such estimates refer to probabilities or variable values that are determined after sample information has been obtained. Posterior estimates essentially involve revising the prior estimates based on new sample information. These topics are beyond the scope of coverage of this chapter, but the interested reader may consult standard statistical textbooks for further information.[8]

D. Interdependencies Among Different Risks

A well thought out risk analysis will consider not only the effects of uncertainty regarding various individual variables, but also how the variables may be interrelated. For example, Table 6-1 indicates an expected research cost of $1 million in Year 1 and an expected development cost of $0.5 million in Year 2. Now suppose these are the only uncertainties that will be investigated. Of course, *actual* costs may be more or less than expected. For the research phase, the actual research cost will be related to the expected cost in the following way:

$$\text{Research Cost} = \text{Expected Research Cost} \times (1 + \text{Forecast Error})$$

(The specific relationship would presumably be defined by specifying a probability distribution.) But it is likely that if research costs are higher than expected, then development costs may be also. Thus, modeling the uncertainty in development costs should attempt to account for the impact of uncertain research costs as well as uncertain development costs. For example, if it were estimated

[8] See, for example, J. McClane and P. Benson, *Statistics for Business and Economics* Dellen Publishing: San Francisco, 1991.

that a 10 percent increase in research costs would lead to an expected 5 percent increase in development costs, then the development cost could be modeled as:

$$\text{Development Cost} = \text{Expected Development Cost} \times (1 + \text{Development Forecast Error} + 0.5 \times \text{Research Forecast Error})$$

It would be relatively easy to incorporate an equation like this into a spreadsheet model, and in fact this is one way Monte Carlo models can be used to specify interdependencies.

In theory, interdependencies among variables can be modeled in a decision analysis framework by explicitly specifying all possible relationships among all variables; of course, this may result in a very complex set of equations even if the underlying benefit-cost flow model is quite small. In practice, the analyst must limit the scope of the analysis by making decisions about which interdependencies to focus on. When using the Monte Carlo approach by employing distributions which allow a range of uncertainty without the need to explicitly specify all possible outcomes, interdependencies can be specified via a "correlation coefficient", which is a number varying between -1 and 1 that indicates the existence or lack thereof of a linear relationship between two variables. A correlation of 1 implies a perfect positive linear relationship. A correlation of 0 implies no linear relationship. A correlation of -1 implies a perfect negative linear relationship.

The reader is referred to standard statistical textbooks for treatment of various probability distributions.[9] In addition, the Monte Carlo simulation software packages discussed below in Section V have certain features to help the user specify interrelationships in a straightforward way.

IV. Sensitivity Testing

In the previous section, it was demonstrated that the outcome of an analysis will depend on numerous factors including estimates of specific variables, forecasts of their future values, assumptions regarding how they vary with each other, and the inherent uncertainties in forecasting future events. Each of these factors has the potential to introduce error into the results. Decision-makers will want to have information on potential errors in order to make informed decisions. In many cases, the degree of uncertainty associated with a particular project may be a decisive factor in determining whether to go forward, either now or in the future, or to cancel the project entirely. This section will demonstrate how to use sensitivity analysis as a way to determine which variables are most likely to have a dramatic effect on a particular project. Thereafter, some of the more sophisticated methods described briefly above will be reviewed.

---

[9] See, for example, W. Mendenhall, *Introduction to Probability and Statistics*, Duxbury: Boston, 1987; J. McClane and P. Benson, ibid.; A. Caniglia, *Statistics for Economics*, New York: HarperCollins, 1992.

The basic approach to sensitivity analysis is to vary key assumptions regarding variables systematically over appropriate ranges and then observe the impact on the net present value of the project. In some cases, the impact may be insignificant, or at least have no affect on the sign of the net present value. For example, in some cases certain variables over a wide range may not alter whether a project has a positive or negative net present value. In such cases, one might conclude that the project itself is insensitive to a particular variable. In other cases, relatively small changes in a particular variable may have dramatic effects on results. Having such information available helps the analyst determine which variables warrant further study. Sensitivity analysis may also be important for a decision-maker in gauging the appropriateness of a project given the willingness of the FAA to accept risk.

Although, as noted above, variations in risk cannot be adequately captured by varying the discount rate, uncertainty as to what is the appropriate discount rate to use in the first place should be captured by evaluating benefit-cost flows at different discount rates. OMB guidance recommends a discount rate of 7 percent.[10] The Office of Aviation Policy and Plans suggests conducting sensitivity estimates at 4 and 10 percent to show the impact of varying the discount rate.

A. One Variable Uncertainty Tests

One useful sensitivity test is to vary one variable at a time, holding all others constant so as to determine the independent, or partial, effect on the outcome. This procedure is known as the one variable uncertainty test. Its primary purpose is to identify the sensitivity of the net present value of each alternative to changes in the value of each component individually.

To carry out the one variable test, the NPV of each alternative must be recalculated for different values of a particular component while all others are held constant. The range of values should extend over those that can reasonably be expected to prevail. Where a probability distribution for a component of interest is known, the relevant range may be established over values corresponding to most of the probability (usually 90 to 95 percent). Where distributions are unknown, the range should extend from the smallest to largest value that could reasonably be expected to occur.

One may then display the results in either tabular or graphic form. In Table 6-1, it was shown that a standard net present value analysis for the ground traffic management technology example showed an expected net present value of $3.5 million. Shown in Table 6-2 are the results of one variable uncertainty tests for the same project. Five variables are varied over appropriate ranges. The five variables are: discount rate, F&E or purchase price of the equipment, operations and maintenance costs (O&M), accident rate, and traffic growth. A quick perusal of the table illustrates that even with relatively wide variations in F&E costs, O&M costs or traffic growth,

10 "OMB Circular A-94" (Revised—October 29, 1992) p. 9.

the ground traffic management technology is estimated to have significant positive net present value. The sole exception relates to changes in the accident rates. In Table 6-1, the technology was assumed to reduce accidents by 50 percent. If this rate is overestimated and in fact there is only a 25 percent reduction, then the one variable uncertainty test suggests that the net present value of the program would be -$640,000. On the other hand, if the accident reduction rate had been underestimated, the net benefits of the technology would be substantially higher than in the base case in Table 6-1. Such an outcome might suggest that the analyst spend additional time assessing the reliability of the accident reduction rate estimates. For example, the analyst might investigate whether the data used to make the estimates were representative of the type of accident that might be prevented using the ground traffic management technology.

TABLE 6-2

ONE VARIABLE TESTS: NPV ($000)

Range	Discount Rate	F&E	O&M	Accident Reduction Rate	Traffic Growth Rate
Lower Value	4,125 (4%)	3,815 (187.5)	4,009 (75)	-640 (25%)	3,503 (0%)
Mean Value	3,498 (7%)	3,498 (250)	3,498 (100)	3,498 (50%)	3,498 (3%)
Upper Value	2,964 (10%)	2,230 (500)	1,455 (200)	7,636 (75%)	3,961 (6%)

B. Two Variable Uncertainty Tests

The one variable test permits examination of one factor holding all others constant. At times it may be useful to let two factors change at the same time. For example, the analyst may be concerned not only about effectiveness of the system, but also about the potential of cost overruns in producing the ground traffic management technology. Concern might focus on what would happen if the accident reduction benefits are overestimated in the base case and the costs of F&E are underestimated. To examine the sensitivity of the project in such an eventuality, the analyst would conduct a two variable test. The results of such a test are illustrated in Table 6-3. The horizontal axis relates to variations in accident reduction effectiveness while the vertical axis relates to variations in F&E costs. The results in the table are the net present values if two variables vary at the levels illustrated on the two axes. All other variables are held constant. For example, the top entry on the left hand side of the table would be a case where F&E costs are actually 25 percent less than the base case (i.e., F&E is $187,500 instead of $250,000 per unit), and where accident reduction effectiveness is 25 percent rather than 50 percent. The result is a net present value of -$323,000. The results from the table also make it clear that the benefits of the project depend much more on the accident reduction estimate than on the F&E estimate.

TABLE 6-3

TWO VARIABLE TESTS: NPV ($000)

F&E ($000)	Accident Reduction Effectiveness		
	25%	50%(Base)	75%
187.5 (-25%)	-323	3,815	7,953
250 (Base)	-640	3,498	7,636
500 (+100%)	-1,908	2,230	6,368

C. Limitations of Sensitivity Analysis

In principle, it would be possible to vary all the variables in an analysis, each over their likely range. One could begin by varying just one variable at a time, then two variables together, then three variables at a time, etc. In this way, sensitivity tests could be conducted on any of the different risk types that typically impact FAA investment projects or regulations. In the context of our example, technological risk could be examined by varying expected accident rate reductions with and without the technology and/or the expected research and development and testing costs; schedule risk could be investigated by changing the year in which different events occur, e.g., the first machine installations at airports might be delayed until Year 4; finally, the impact of external uncertainty could be examined by changing the assumed growth rate of airport operations.

While the potential ramifications of all sorts of changes could be investigated in this way, such an approach would quickly yield an overload of information. A practical goal of risk analysis is to present a range of likely results in a compact, straightforward manner, and presenting a large number of sensitivity results is not likely to be a satisfactory way to achieve this goal.

Moreover, sensitivity testing does not require one to assess how likely it is that specific values of the parameters at issue will actually occur. In the above one-variable tests for accident reduction effectiveness, one cannot adequately assess how uncertainty affects net present values without some knowledge about the likelihood that accident reduction effectiveness will actually be 50% or 100% above or below the base expected value.

Sensitivity testing also does not encourage the analyst to consider relationships *between* parameters and the probabilities that certain values will occur together. In the two-variable tests discussed above, for example, it may well be that higher-than-expected F&E expenditures may result in higher-than-expected accident reduction effectiveness, and vice versa. In such a situation, NPV results in Table 6-3 from the lower right corner and upper left corner are more relevant than those in the lower left corner or upper right corner, yet this is not reflected in the table itself.

These criticisms of sensitivity testing can, in many cases, be overcome by the careful structuring of a set of sensitivity scenarios that accounts for the distribution of parameter probabilities and interrelationships among the variables being tested. In such an analysis, interdependent restrictions would be placed on the variables and they would be varied as a group with the restrictions in place; each such set of restrictions would constitute a scenario. With such an approach, it is important to ensure that each scenario is internally consistent (i.e., that the implied or explicit interdependencies make sense). In this way it is possible to create a set of scenarios which are representative of the universe of possible outcomes.

A more formal approach to risk analysis which forces the analyst to directly consider both the probability distribution of specific parameters (i.e., how likely is a given value of a parameter), as well as the interrelationship between values of different parameters, is offered by Monte Carlo analysis.

V. MONTE CARLO ANALYSIS

In the preceding sections, it has been shown that sensitivity analysis can be useful in considering the effects of uncertainty regarding critical variables in a capital budgeting exercise. By looking at a project under alternative scenarios, one can consider the effect of a limited number of plausible combinations of variables. Monte Carlo simulation is a tool for considering many more possible combinations. It uses simulation techniques to calculate the entire range of all possible project outcomes *and* the likelihood of each actually occurring. When collected in a probability or frequency distribution, this information presents decisionmakers with a concise summary of a project's benefit-cost status.

To under take a Monte Carlo analysis, the analyst must be able to specify the determinants of each of the variables under consideration, and their range as expressed in probability distributions, and the interrelationships between them. For most real world applications, it is impossible to fully specify every variable, all of their determinants and/or their interrelationships. But, using computer modeling packages designed for the purpose, an analyst can select the distribution of a particular variable. This process allows the analyst to skip the formal specification of the determinants of each variable in exact detail, and instead rely upon generalized distribution forms which are likely to be representative. For example, accident rates might be best represented by Poisson distributions, while certain other technical variables like installation costs might vary according to the normal distribution. Simulation packages are designed to help the analyst select both the shape of the distribution and the size of the variation in each distribution. The packages can also account for interdependencies between variables as specified by the analyst. The computer then models the results by randomly selecting values for each variable from the applicable distribution and then computing NPV's for the project.

Under the Monte Carlo approach, hundreds or thousands of simulations of the model are performed. For each simulation, the computer randomly samples from the probability distribution for each variable (accounting for interdependence if necessary), and computes and stores the results. In the present context, the final result is a distribution of net present values for the project. Some software packages can provide additional information regarding which variables have the most influence on the present value distribution, the likelihood of achieving a target net present value, etc.

A. Conducting a Monte Carlo Simulation

As mentioned earlier, the Monte Carlo method is particularly well-suited for conducting risk analysis when it is not feasible to explicitly identify and list all possible discrete outcomes and their likelihood of occurrence. This is most often the case when trying to assess uncertainty in safety or accident rates, cost uncertainty with respect to technical, R&D, or operations costs, or the risk associated with uncertain estimates of recurring future activities such as growth rates in airport operations, inflation rates, etc. In these types of situations, it is most often reasonable to specify uncertainty via a probability distribution which can be specified with only a few parameters (e.g., the mean and variance), yet that adequately captures the range and likelihood of uncertain outcomes. Monte Carlo analysis is designed specifically for such situations.

Consider again the ground traffic management technology analysis from Table 6-1. Suppose we are interested in assessing how the results are affected by uncertainty with respect to the cost of F&E (currently estimated at $250,000 per machine) and the accident rate reduction actually achievable by the technology (currently estimated at 50 percent). Note that the associated NPV results are affected by uncertainty on both the cost side (due to F&E) and the benefit side (due to the reduced accident probability).

To illustrate how a simple Monte Carlo exercise could be undertaken, assume that the F&E estimate is subject to some uncertainty; in particular, we believe that this uncertainty can be adequately captured by specifying the F&E cost per machine as a normal random variable with a mean of $250,000 and a standard deviation of $50,000. Remembering that about 95 percent of the probability distribution lies within two standard deviations of the mean, this implies a 95 percent chance that the true cost will lie in the range of $150,000 to $350,000. In turn, this means that when the computer is randomly drawing values for the F&E cost, about 95 percent of them will be within this range.

At the same time, let us assume that the accident reduction rate, currently set at 50 percent, involves uncertainty that can be adequately characterized by a triangular distribution, with a most likely value of 70 percent, a minimum of zero percent (implying no impact on the accident rate), and a maximum of 80 percent. Such a distribution implies a *mean* value of 50%, consistent with the base expected value used in Table 6-1. Finally, let us allow for interdependence between these two variables. Recalling the earlier discussion that higher-than-expected F&E costs are likely to lead to higher-than-expected accident rate reductions, this implies a *positive* correlation between the two variables; assume that we estimate this correlation to be 0.5. With this specification, the computer is being instructed to ensure that the majority of the simulations involve either a high F&E coupled with a high accident reduction rate, or a low F&E coupled with a low accident reduction rate; there will be relatively few simulations where a high F&E occurs along with a low accident reduction rate, or vice versa.

FIGURE 6-1

MONTE CARLO SIMULATION RESULTS for F&E and ACCIDENT RATE REDUCTION VARIATIONS PRESENT VALUE of NET BENEFIT ($ 000)

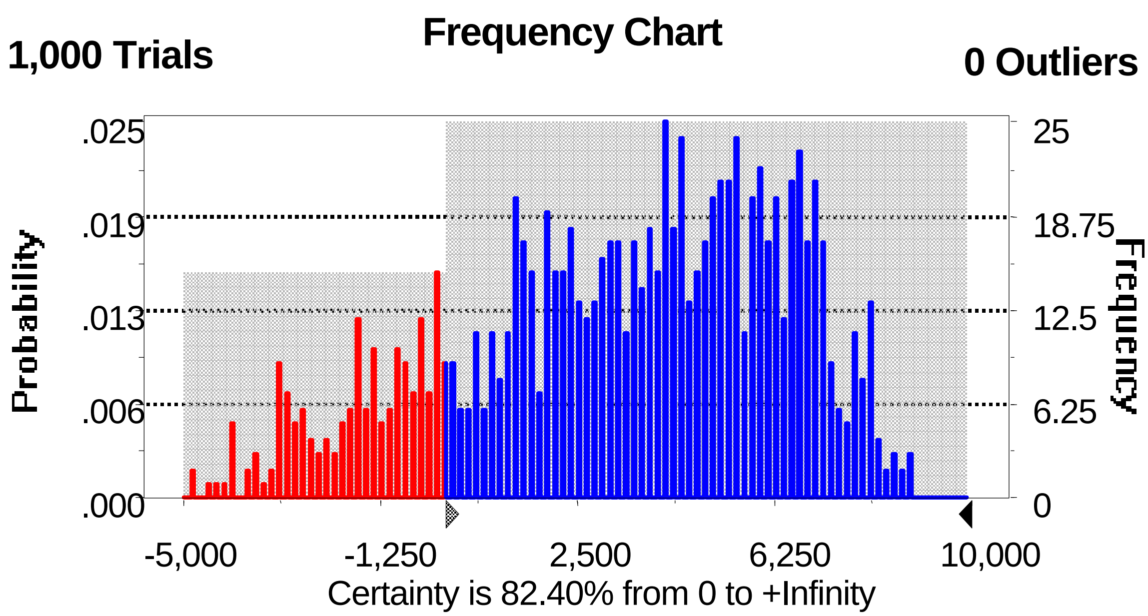

The NPV results of a Monte Carlo analysis with these specifications, running 1,000 simulations, are shown in Figure 6-1. Note that the majority of the distribution represents positive values, implying a net return of 7 percent or greater. In fact, the simulation results indicated that 82.4 percent of the distribution of the present value of net benefit was in the positive range. Thus, the analyst could conclude, based on this set of simulations, that there is about a 4 in 5 chance that the project would generate a positive net present value using a discount rate of 7 percent. To assess whether enough simulations have been run to ensure accurate results, one can compare the "observed" mean NPV that occurred when both the F&E cost and the accident rate were set to their mean values ($3,498,000 from Table 6-1) with the mean NPV from the simulations ($3,352,000). This indicates that the simulation results are consistent with the mean results from the initial spreadsheet analysis. Other useful outputs can be obtained from a Monte Carlo analysis, but the NPV distribution is the single most key result to be utilized by the analyst.

B. Using Commercially Available Monte Carlo Software

At the time of this writing, there are at least two commercially available computer packages that work in conjunction with either Lotus 1-2-3© or Microsoft Excel© spreadsheets and produce Monte Carlo results.[11] Both of these programs start with a spreadsheet representation of a project to be analyzed such as those shown in the previous tables. The Monte Carlo programs then work directly in conjunction with the spreadsheets to produce the simulation results.

The process is relatively straightforward. The analyst identifies the variables within the spreadsheet that are subject to uncertainty. The distribution that best characterizes the range of possible values is then selected for each variable. For example, the analyst might select a normal distribution for certain variables or Poisson distribution for others. The model would then prompt the analyst for additional information that would allow it to apply the distribution to that variable. As was shown above, the analyst can also make one variable depend on another. For example, one might calculate Variable B based upon the value of Variable A. In effect this would mean that values selected by the computer from the distribution for Value A would determine (in part) the value of Variable B.

The simulation packages are designed to accommodate virtually any type of statistical distribution for any type of variable. They then use random number generators to select values from the distributions for each variable and assign them for a particular simulation. Because oftentimes more than one variable is subject to uncertainty, the possible combinations of outputs becomes very large. One advantage of the software packages is that they take advantage of the capability of a computer to make hundreds or thousands or iterations of the model in a relatively short period of time. The packages then produce all of the tabular and graphical representations of risk presented in this chapter, as well as many more. For example, the packages quickly produce a rank ordering of the sensitivity of models to each individual variable. Or, one can quickly identify which variables account for relatively high or low project outcomes.

Another key advantage of these risk analysis packages is that they are self documenting. That is, one can quickly access a list of assumptions made in the risk analysis model, and alter them as appropriate as new information becomes available. It is therefore recommended that an analyst take advantage of commercially available software, especially when complicated investment decisions are to be considered.

---

[11] The two software products are *@Risk*© (Palisade Corporation, 31 Decker Road, Newfield, NY 14867, www.palisade.com) and *Crystal Ball*© (Decisioneering, Inc., 1515 Arapahoe Street, Suite 1311, Denver, CO 80202, www.decisioneering.com)

C. Limitations of Monte Carlo Analysis

Monte Carlo simulation forces decision-makers to explicitly consider uncertainty and interdependencies among different inputs, and can be a useful tool in assessing risk. On the other hand, it is often difficult to examine certain types of risks associated with schedule uncertainty in a Monte Carlo framework. On a theoretical level, this is because one may not have any reasonable basis on which to ascribe a parameterized probability distribution to the schedule of relevant events, e.g., R&D may be expected to take two years, but it may be difficult to assess, say, the variance of this estimate. Even if one could make such an assessment, however, the standard spreadsheet framework for designing Monte Carlo experiments does not lend itself to incorporating schedule risk except in a very simple way. As our earlier example shows, a typical spreadsheet analysis "hard-wires" the schedule for when events are expected to occur, e.g., development and testing begins in Year 2. One way to account for uncertainty in this estimate would be to create a second testing phase that begins in, say, Year 3, and set up the model so that the Monte Carlo technique randomly select only one of the two possible testing phases for each simulation. Of course, another alternative would be to carry out sensitivity tests and scenario analyses out by manually adjusting the schedule between successive analyses.[12]

A second drawback of Monte Carlo simulation comes when attempting to interpret its primary output, i.e., a probability distribution of net present values. The amount of risk is presumably reflected in the dispersion of the NPV distribution. But it is sensitive to the definition of the project being analyzed. For example, if two unrelated projects were combined and analyzed as one, the "risk" of the NPV of the combined project would be less than the average "risk" of the NPV's of the two separate projects. Moreover, it is difficult to interpret a distribution of NPV's. There is no single rule arising out of such a distribution that can guide decision-makers concerning the acceptable balance between expected return and the variance of that return.

VI. Decision Analysis

Another way that an analyst can help improve decision-making regarding investment projects is to recognize that many investment evaluations involve not one decision but rather several sequential decisions. If subsequent investment decisions depend on those made today, then today's decision may depend on what happens tomorrow. By recognizing that decisionmakers may be able to stop projects after one or two steps in the process without committing the full measure of resources,

---

[12] If one were not tied to a spreadsheet format, it is possible that one could handle schedule uncertainty via Monte Carlo methods. A simulation technique known as VERT (Venture Evaluation and Review Technique) has been designed which allows one to analyze schedule and time-based risks. This technique has been used in previous FAA-sponsored analyses of advanced automated air traffic control systems. See *Quantitative Assessment of Risks on the Attainment of the Benefits and Costs of Advanced Automated Air Traffic Control System Alternatives, Volume VI*, prepared by Mitre Corporation, 1987.

the analyst may provide options for decision-making that would otherwise go unrecognized. These concepts are best illustrated using decision trees.[13]

To illustrate the value of decision trees, Figure 6-2 illustrates many of the important features of the example for the ground traffic management technology first summarized in Table 6-1. For purposes of this example, the earlier assumptions that the initial research and subsequent development and testing will necessarily be successful are changed to allow for the possibility of failure. Specifically identified are decision points (designated by squares) and chance nodes (designated by solid dots). In Figure 6-2, there are three distinct decision points that have been identified by the analyst (although there could easily be many more depending on which outcomes in the process of researching, testing and implementing the ground traffic management technology are actually uncertain). The three key decision points are:

1. Initiate research on the ground traffic management technology: This is followed by a chance node which can result in success or failure. Based on information collected from technical personnel, the analyst assigns a probability of success of 50 percent and a probability of failure of 50 percent.

2. Develop and test the equipment: This is followed by a chance node which can result in success or failure. The analyst estimates the probability of success at 75 percent and a probability of failure at 25 percent.

3. Implement equipment: The analyst assumes implementation will result in the full measure of benefits assumed in Table 6-1. (This is, the chance node has a probability of success of 100 percent.)

[13] Many of the same concepts are used in option trading in the stock market although the presentations are much more difficult to follow and are not generally relevant to project decision-making.

Figure 6-2
Decision Tree Analysis of Project
($000)

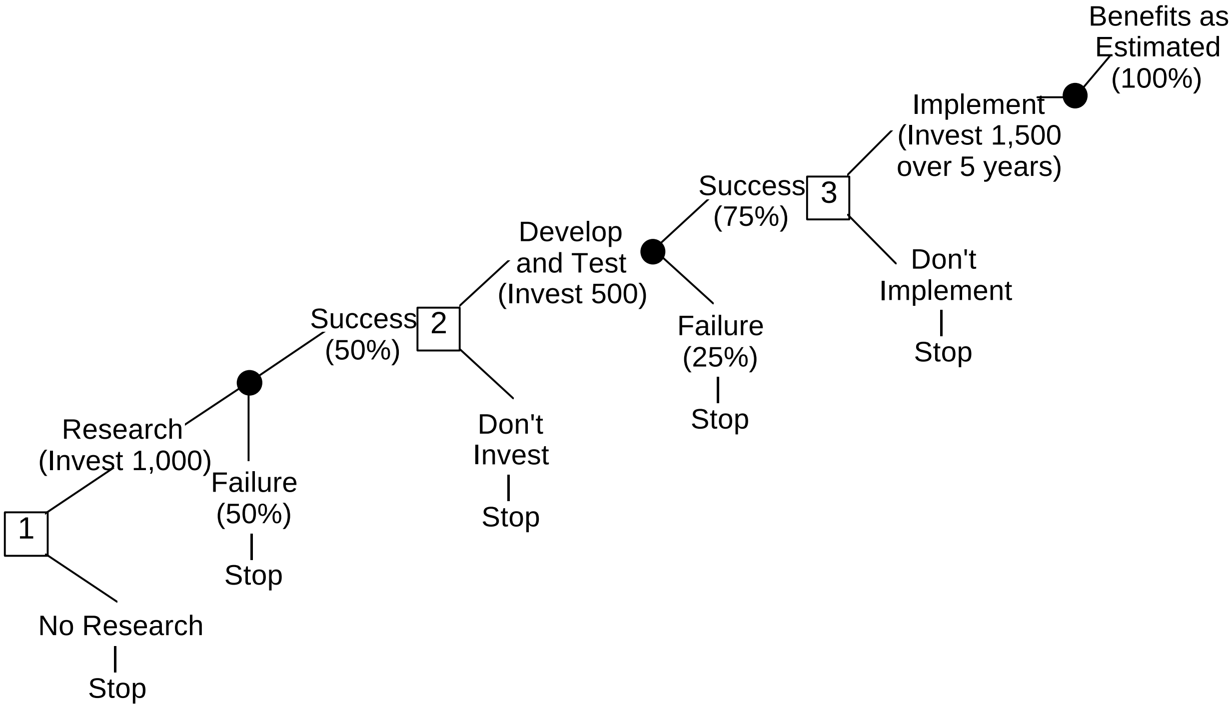

As Figure 6-2 shows, at each decision point, the analyst has allowed for the possibility of stopping the process. That is, the decisionmaker might decide not to do any research in which case the whole program would be canceled; alternatively, the decisionmaker may have the option of canceling the program after the research phase, or after the development and testing phase. As will be shown below, there is sometimes significant value in having these or other types of options.

The main problem confronted by the decisionmaker is whether to go forward with the research program today. To solve the problem, one begins first with what to do if confronted with the implementation decision; that is, one begins at the right side of the decision tree. Table 6-4 illustrates the implementation decision. The only decision that would have to be made in the third year of the project is whether to implement it. Table 6-4 shows that the net present value of implementation is positive, when expressed in Year 3 dollars. That is, standing in Year 3, the analyst would conclude that going forward with the implementation of the ground traffic management project would have a significant positive benefit.

TABLE 6-4

IMPLEMENTATION DECISION ($000)

Year	1	2	3	4	5	6	7
Investment			-1,000	-250	-250	0	
Operating Costs			-400	-500	-600	-600	-600
Benefits			2,042	2,103	2,167	2,232	2,299
NPV (Year 3)			642	$\frac{1,353}{1.07}$	$\frac{1,317}{(1.07)^2}$	$\frac{1,632}{(1.07)^3}$	$\frac{1,699}{(1.07)^4}$
Total NPV (Year 3)			5,685				

The only decision that would have to be made in Year 2 is whether to go forward with the development and testing phase of the program. The analysis of this decision is illustrated in Table 6-5. It uses results from the implementation decision together with other information from the decision tree. We know that the probability of success of development and testing is 75 percent; we also know from Table 6-4 that if development and testing is successful and the decision is made to go forward, then the project will realize a benefit of $5.685 million. If development and testing fails, however, there will be zero benefits because the program will be canceled. In Table 6-5, the expected values for success are calculated as the sum of:

$$[(75\%) \times \$5,685] + [(25\%) \times 0] = \$4,264.$$

That is, the expected value of the benefits from development and testing will be realized in the implementation phase and would have a value in Year 3 of $4.264 million, once the probabilities of success and failure are taken into account. This value is then discounted back into Year 2 dollars and added to the investment costs of development and testing. As shown in Table 6-5, the total net present value of development and testing is also quite positive — $3.485 million. So far, the decision tree analysis suggests that both implementation, and development and testing are likely to have strongly positive net benefits.

TABLE 6-5

DEVELOPMENT and TESTING DECISIONS ($000)

Year	2	3
Investment	-500	
Benefit: Success (75%) Failure (25%)	 -500	 5,685 0 4,264
NPV (Year 2)	-500	4,264 1.07
Total NPV (Year 2)	3,485	

There is now enough information to examine the research decision. Recall from the decision tree that the investment cost of the research program is $1 million and that it has a probability of success of 50 percent and an equal probability of failure. From Table 6-5 we already know that the net present value of the program if the research effort is successful is $3.485 million. If the research program fails, the payoff is zero. Using the same expected value techniques described above, we find, as shown in Table 6-6, that the expected benefits from the program in Year 2 dollars is $1.743 million, which when discounted back to Year 1 and added to the investment costs shows a total net present value for the research program of $629,000. Thus, the ground traffic management technology continues to have an expected positive net present value under assumptions which allow for the chance of failure in the research and development and testing phases.

Once the probabilities are taken into account, the analyst could then repeat the one variable and two variable sensitivity tests described previously.

TABLE 6-6

RESEARCH DECISION ($000)

Year	1	2
Investment	-1,000	
Benefit: Success (50%) Failure (50%)		 3,485 0
NPV (Year 1)	-1,000	1,743 1.07
Total NPV	629	

A. Irreversibility and Abandonment

One of the important issues faced by the FAA is that many of its investment programs are irreversible. That is, once money is spent on a project, it is difficult to recoup any money from it should that program be discontinued because many of the investments made by the FAA tend to be in specialized equipment as opposed to generally available equipment. For example, in industry an airline making an investment in an aircraft does not face an irreversible decision since it is likely that there will be a secondary market in the asset should the airline decide to abandon its use. This may not be the case for all FAA programs if much of the equipment has no other alternative use because it was developed for a specialized purpose like air traffic control.

There is sometimes a significant value in having the option to abandon an investment, and this can be illustrated using the decision tree process described immediately above. Suppose for the moment that part of the research program for the ground traffic management project would involve purchasing general purpose computer equipment which could be sold a year later for $200. That is, part of the $1 million investment in the research program could be reversed if the research program were to fail. Table 6-7 illustrates a method for valuing having this option. The analysis is identical to the one shown in Table 6-6 with the exception that in the event the research program fails the FAA is able to realize $200,000 from the sale of the computer equipment. This raises the expected payoff from the program resulting in a net present value of $722,000.

TABLE 6-7

VALUE of ABANDONMENT OPTION after FAILURE of RESEARCH ($000)

Year	1	2
Investment	-1,000	
Benefit: Success (50%) Fail and Abandon (50%)		 3,485 200
NPV (Year 1)	-1,000	1,843 1.07
Total NPV	722	

Value of Abandonment Option = NPV with abandonment minus NPV without abandonment
= 722 - 629
= 93

The value of being able to abandon the program and reverse some of the investment is merely the difference between the net present value with abandonment and the net present value without it, or $722,000 - $629,000. That is, it is worth $93,000 to the FAA to have the option of selling the computer equipment, or when the FAA is developing a contract that may have an escape clause that allows it to discontinue a program. Having this kind of information may be extremely important to decisionmakers when there are alternative methods for conducting programs with general purpose versus specialized equipment.

B. Valuing Reductions in Uncertainty

Suppose that some of the uncertainty regarding the ground traffic management project could be eliminated if additional research funds were allocated to the project. One can use the decision tree analysis to calculate the value of this uncertainty reduction in order to improve the reliability of the variables included in the analysis.

For example, let us suppose that the 50 percent estimated success rate for the research phase is deemed to be too low. This might happen because the analyst is concerned with the reliability of the 50 percent estimate, or because the net present value estimate for the project is too low. In either case, paying now to reduce uncertainty tomorrow might be beneficial. To determine this, one could posit a circumstance where additional information could improve the chance of success for the development and testing phase of the program. Let us suppose we are interested in knowing how much it is worth to eliminate any uncertainty about the research phase of the program — i.e., that the success rate would increase to 100 percent. How much is that worth? Because the research success rate is now 100 percent, the expected value increases. In order to gain that increase in reliability, suppose that an additional $1 million must be committed in Year 1. The value of the additional research expenditures are computed in Table 6-8.

TABLE 6-8

VALUE of REDUCTION in UNCERTAINTY ($000)

Year	1	2
Investment	-2,000	
Benefit: Success (100%) Failure (0%)		 3,485 0
NPV (Year 1)	-2,000	$\frac{3,485}{1.07}$
Total NPV	1,257	

Value of Uncertainty Reduction = 1,257 - 629
= 628

The result is that the net present value of the beefed-up research program with a 100 percent success rate would be $1,257,000 as compared with the estimated $629,000 in Table 6-6. The value of the option of spending additional funds to eliminate research uncertainty is therefore worth $628,000 — a very substantial percentage of the estimated base case. It appears that improving the success rate by spending additional research funds would be of substantial value to the FAA and should be pursued.

This sort of analysis can also be applied to other sorts of options that may be available, e.g., the option to spread out expenditures over a longer period of time, the option to delay the testing and development phase until Year 3, etc.

C. Limitations of Decision Analysis

Decision analysis can be a valuable tool in assessing the risks associated with potential projects that unfold over time because it allows decision-makers to assess projects at certain key points in the process without committing the full measure of financial or other resources. But the requirements needed to conduct a useful analysis can be difficult to meet. The most important practical requirement is that the uncertainty of all relevant variables must be characterized in terms of well-defined alternative discrete events. As described earlier, this may or may not be achievable depending on the sorts of risks being considered. In practice, schedule risks and accident risks may be more amenable to the assigning of probabilities and explicit outcomes to each possible event than other risks associated with, say, technical or compliance risks. Careful judgment is required in assessing whether a decision approach to risk assessment is warranted for any particular application.